Games in extensive form - basic concepts and definitions

Wiesław Zielonka

www.liafa.univ-paris-diderot.fr/~zielonka

mail: zielonka@liafa.univ-paris-diderot.fr

LIAFA, Université Denis Diderot

1 Perfect information games

A perfect information game in extensive form is composed of the following elements:

- a finite set $N = \{1, \ldots, n\}$ of players,
- a non-empty set V of positions with a partition $V = V_0 \cup V_1 \cup \ldots \cup V_n \cup T$, where V_i, $i \in N$ is the set of positions controlled by player i, V_0 is the set of chance positions and T is the set of terminal positions,
- for each non-terminal position $v \in V \setminus T$ a finite non-empty set $A(v)$ of actions available at v,
- a transition mapping δ which for each non-terminal position $v \in V \setminus T$ and each action $a \in A(v)$ gives a position $\delta(v, a) \in V$ reachable from v upon the execution of a,
- a mapping p which for each chance position $v \in V_0$ gives a probability distribution over the set $A(v)$ of action available at v. We write $p(a|v)$ to denote the probability of action $a \in A(v)$ at $v \in V_0$ and we add the requirement that $p(a|v)$ is always strictly greater than 0 (we can always remove from $A(v)$ the actions that are taken with probability 0),
- a family $u_i, i \in N$ of payoff or utility mappings, $u_i : T \cup \text{Paths}^\omega \longrightarrow \mathbb{R}$, where for each terminal position $v \in T$, $u_i(v)$ gives the payoff of player i if the game terminates at v and for each infinite path $p \in \text{Paths}^\omega$, $u_i(p)$ is the payoff player i receives for an infinite play p.

We suppose that all this structure satisfies one additional condition. Let

$$E = \{(v, w) \in V \times V \mid v \in V \setminus T \text{ and } w = \delta(v, a) \text{ for an action } a \in V(v)\}$$

then (V, E) is an oriented tree, the root of this tree is the initial game position and terminal positions T correspond to the leafs and Paths^ω is the set of infinite paths in (V, E) starting at the root.

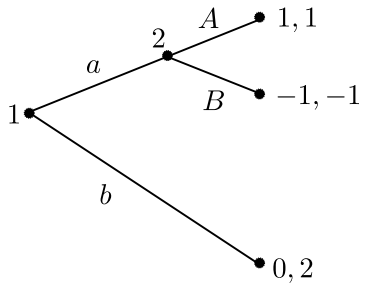

Figure 1: An extensive game with two players and no chance positions. Vertices are labeled with players' names. The initial position is controlled by player 1 who can execute either action a or action b. The execution of b leads to a terminal position with payoff 0 for player 1 and payoff 2 for player 2. The execution of action a leads to a position controlled by player 2, where he can execute either action A or action B.

A *pure* strategy σ_i for player $i \in N$ in a perfect information extensive game is a mapping that for each position $v \in V_i$ gives an action $\sigma_i(v) \in A(v)$ that player i executes at position v.
A *mixed* strategy for player $i \in N$ is a probability distribution over pure strategies of player i.
The intuition behind mixed strategies is that player i uses randomization at the beginning of the game to choose one of his pure strategies and subsequently he plays using the chosen pure strategy.
A *behavioral* strategy σ_i for player $i \in N$ is a mapping that for each position $v \in V_i$ controlled by player i gives a probability distribution $\sigma_i(v) \in \Delta(A(v))$ over actions $A(v)$ available at v. As usually we shall note by $\sigma_i(a|v)$ the probability that strategy σ_i assigns to action a at v.
A *play* in the game is a path in the game tree from the root to a leaf, i.e. it is a sequence $v_0a_0v_1a_1 \dots v_ka_kv_{k+1}$ such that v_0 is the initial game position (the root of the game tree), v_{k+1} is a terminal position (a leaf in the game tree) and for each l, $1 \leq l \leq k$, $v_{i+1} = \delta(v_i, a_i)$.
If the set of positions V is finite then a given profile of behavioral strategies $\sigma = (\sigma_1, \dots, \sigma_n)$ yields a probability distribution over terminal positions. Let $t \in T$ be a terminal position and r the initial position of the game. Let $v_0a_0v_1a_1 \dots v_ka_kv_{k+1}$, $v_0 = r$, $v_{k+1} = t$ be the play from r to t. Then the probability $P(t|\sigma)$ that t is reached under profile σ is given by multiplying the probabilities along this play: $P(v_0, a_0)P(v_1, a_1) \dots P(v_k, a_k)$, where $P(v_l, a_l) = p(a_l|v_l)$ if $v_l \in V_0$ is a chance position and $P(v_l, a_l) = \sigma_i(a_l|v_l)$ is the probability assigned by strategy σ_i to action a_l at position v_l if $v_l \in V_i$ is a position controlled by player i.
The payoff of player $i \in N$ for a strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ is defined through expectation:

$$u_i(\sigma) = \sum_{t \in T} u_i(t) \cdot P(t|\sigma). \tag{1}$$

Since pure strategies can be seen as a special kind of behavioral strategies that assign only probabilities 0 and 1 to actions available at each vertex the definition above applies also to define payoffs for pure strategy profiles.

If we have payoffs for pure strategy profiles then, taking expectation, we get also the payoff for mixed strategy profiles.

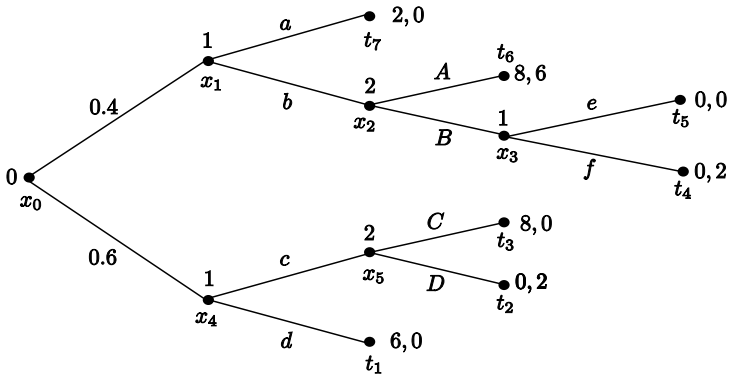

Figure 2: The unique chance position is x_0 from which with probability 0.4 the game moves to x_1 and with probability 0.6 to x_4.

Example 1. In the game on Figure 2 there is one chance move at the beginning. Player 1 has 8 pure strategies: ace, acf, ade, adf, bce, bcf, bde, bdf. Player 2 has 4 pure strategies: AC, AD, BC, BD (since we have chosen unique actions there is no need to specify the positions).
A possible mixed strategy for player 1 is $\frac{5}{15}ace+\frac{3}{15}adf+\frac{7}{15}bde$, where the coefficients indicate the probabilities of playing one of the three pure strategies.
A possible behavioral strategy for player 1 is $\sigma(\cdot|x_1)=\frac{1}{2}a+\frac{1}{2}b$, $\sigma(\cdot|x_3)=\frac{2}{3}e+\frac{1}{3}f$, $\sigma(\cdot|x_4)=\frac{1}{5}c+\frac{4}{5}d$, which means that for example at x_3 player 1 chooses e with probability 2/3 and f with probability 1/3.

1.1 Infinite perfect information games without chance positions

Let A be a finite alphabet, A^* the set of all finite words over A, A^ω the set of all infinite words over A, i.e. the set of all infinite sequences $a_1a_2a_3\ldots$ with $a_i\in A$.
We take A^* as the set of positions of a two-player game. For each $u\in A^*$, executing action $a\in A$ moves us to a new position ua. The set of terminal nodes is empty. We assume that player I controls positions $\bigcup_{i=0}A^{2i}$ of even length while player II controls positions $\bigcup_{i=0}A^{2i+1}$ of odd length. Thus starting at position ε the players choose actions alternately, player I choose a_1, player II a_2, player I a_3 etc. Let $W_I\subset A^\omega$ be the set of infinite words winning for player I, i.e. player I wins if the infinite word constructed during the play belongs to W_I, $W_{II}=A^\omega\setminus W_I$ are infinite words winning for player II. Thus this is a 0-sum game when player I receives payoff 1 if the resulting infinite word is in W_I and he receives payoff -1 if the resulting infinite word is in W_{II}.
We can define a metric d on A^ω, for two infinite words $u=a_1a_2\ldots$, $v=b_1b_2\ldots$ we set

$$d(u,v)=\begin{cases}0 & \text{if } u=v,\\ 2^{-n+1} & \text{if } a_n\neq b_n \text{ and } a_i=b_i \text{ for all } i<n\end{cases}$$

In the topology induced by this metric the open sets have the form LA^ω, where $L \subset A^*$.
We say that the game with the winning set $W_I \subset A^\omega$ is determined if one of the players has a winning strategy (i.e. if the games has a value).

Theorem 1 (Gale and Stewart, 1953)**.** *If W_I is open or closed then the game is determined.*

Proof. Suppose that W_I is open, i.e. $W_I = LA^\omega$ for some $L \subseteq A^*$. The set A^* of finite words is the set of positions in the game. The aim is to find the set of winning P_I positions for player I. Intuitively, P_I is the set of positions such that player I has a winning strategy for the games starting at $x \in P_I$. Thus if the empty word ε belongs to P_I then player I has a winning strategy for the game starting with the empty word.
We write $x \leq y$ to denote that a finite word x is a prefix of y, i.e. that there exists $u \in A^*$ such that $xu = y$.
We define by induction a sequence of position sets $(P_i)_{i\geq 0}$.
We start with

$$P_0 = \{y \in A^* \mid x \leq y \text{ for some } x \in L\}$$

and next

$$\begin{aligned} P_{i+1} = P_i \cup \\ \{x \in A^* \mid x \text{ has an even length and there exists } a \in A \text{ such that } xa \in P_i \} \cup \\ \{x \in A^* \mid x \text{ has an odd length and, for each } a \in A,\ xa \in P_i \}. \end{aligned}$$

We set

$$P_I := \bigcup_{i_0}^{\infty} P_i$$

and we claim that player I has a winning strategy for games starting in P_I while player II has a winning strategy for games starting in $P_{II} = A^* \setminus P_I$
The proof is again by induction.
If $x \in P_0$ then position x is clearly winning for player I since all infinite words with prefix x belong to W_I and player I can play any action at each stage and the resulting infinite word will belong to W_I whatever the moves of player II.
Assume that all initial positions in P_i are winning for player I and let $x \in P_{i+1}$. If x has even length then x is a position where player I plays. From definition of P_{i+1} is follows that there exists an action $a \in A$ such that $xa \in P_i$. Thus w winning strategy of player I will be :
play an a at x such $xa \in P_i$ and use his winning strategy in P_i in the sequel.
Clearly with this strategy player I wins for games starting in P_{i+1}.
It remains to prove that player II has a winning strategy for games starting in positions belonging to $P_{II} = A^* \setminus P_I$.
If $x \in P_{II}$ is of even length (a position controlled by player I) then for each action $a \in A$, xa belongs to P_{II}. (Indeed suppose that $xa \in P_I$. Then $xa \in P_i$ for some i implying that $x \in P_{i+1} \subseteq P_I$.)
If $x \in P_{II}$ is of odd length (a position controlled by player II) then there exists an action a such that $xa \in P_{II}$. (Indeed if for all actions a, $xa \in P_I$ then for each action

a there exists an integer i_a such that $xa \in P_{i_a}$. However there are only finitely many actions thus $j = \max\{i_a \mid a \in A\}$ is well defined and we can see that $xa \in P_j$ for all a. But this would imply that $x \in P_{j+1} \subseteq P_I$.)
Thus the strategy of player II on P_{II} is to choose at each $x \in P_{II}$ of odd length an action a such that $xa \in P_{II}$. Clearly if player II plays using this strategy then the game starting in any position of P_{II} will remain in P_{II} and will never hit P_I and, therefore, it will never visit the positions of L. Thus the constructed infinite word will not belong to W_I.
If W_I is a closed set the the proof is similar since W_{II} is open (just exchange players I and II). □

In the proof above we used (where?) the fact that A is finite. What about infinite A? The proof idea is still valid but this time we should construct P_I by a transfinite induction.
Le $\mathcal{B}$ be the family of Borel sets over A^ω. This is the smallest family containing all open sets and closed under countable unions, countable intersections and complement.

Theorem 2 (Martin)**.** *If $W_I \in \mathcal{B}$ then the game is determined.*

For the proof you can consult the monograph of Kechris, Classical Descriptive Set Theory, Springer. The result of Martin is one of the cornestones of game theory.

Theorem 3 (Gale and Stewart)**.** *There exists $W_I \subset A^\omega$ such that the game with the winning set W_I for player I is not determined.*

Proof. Of course such a set W_I cannot be Borel, we cannot construct W_I starting from open sets and using countable unions, countable intersections and complements, the set W_I such that the game is not determined is very complicated.
To understand the proof it is better to know what is an ordinal number and what is a cardinal.
Let $A = \{0, 1\}$ and let $V_I = \bigcup_i A^{2i}$, $V_{II} = \bigcup_i A^{2i+1}$ be the sets of positions controlled by I and II respectively. Thus a strategy of player I is a mapping $\sigma : V_I \to \{0, 1\}$ while a strategy for player II is a mapping $\tau : V_{II} \to \{0, 1\}$.
For each strategy σ of player I let $H(\sigma)$ be the set of all infinite words $u \in A^\omega$ that are compatible with σ. By definition $H(\sigma)$ consists of all infinite words u such that for each finite prefix $v \in V_I$ of u the word $v\sigma(v)$ is a prefix of u.
The set $H(\tau)$ of infinite words consistent with a strategy τ of player II is defined in a similar way.
If player I plays using a strategy σ while player II plays using τ then they construct a unique infinite word u, in fact this is a unique word in the intersection $H(\sigma) \cap H(\tau)$.
Strategy σ is winning for player I if player I wins against any possible strategy of player II. However this is possible if and only if $H(\sigma) \subseteq W_I$.
Since V_I is countable the cardinality of the set Σ of strategies of player I is $2^{\aleph_0}$, where $\aleph_0$ is the first infinite cardinal.
Let α be the least ordinal number such that the set $\{\beta : \beta < \alpha\}$ of ordinals smallest than α has cardinality $2^{\aleph_0}$.
Thus we can index all strategies of players I and II with ordinals smaller than α, i.e. we can assume that $\{\sigma_\beta, \beta < \alpha\}$ and $\{\tau_\beta, \beta < \alpha\}$ are the sets of all strategies of players I and II respectively.

Note that the cardinality of each set $H(\sigma)$ is $2^{\aleph_0}$ (there is a natural bijection between $H(\sigma)$ and the set A^ω of all infinite words).
Now we define by transfinite induction two sequences (indexed by ordinals) of infinite words:

- choose $y_0 \in H(\tau_0)$ and then choose $x_0 \in H(\sigma_0)$ different from y_0.
- suppose that x_γ, y_γ are defined for $\gamma < \beta < \alpha$, then choose $y_\beta \in H(\tau_\beta) \setminus \{x_\gamma \mid \gamma < \beta\}$ and $x_\beta \in H(\sigma_\beta) \setminus \{y_\gamma \mid \gamma \leq \beta\}$. This step is always possible since $H(\tau_\beta)$ and $H(\sigma_\beta)$ have cardinality $2^{\aleph_0}$ while the cardinality of $\{x_\gamma \mid \gamma < \beta\}$ and $\{y_\gamma \mid \gamma \leq \beta\}$ is smaller than $2^{\aleph_0}$ implying that $H(\tau_\beta) \setminus \{x_\gamma \mid \gamma < \beta\}$ and $H(\sigma_\beta) \setminus \{y_\gamma \mid \gamma \leq \beta\}$ are nonempty.

The sets $X = \{x_\beta \mid \beta < \alpha\}$ and $Y = \{y_\beta \mid \beta < \alpha\}$ are disjoint. Indeed assume that $x_\beta = y_\gamma$ for some $\beta, \gamma < \alpha$. If $\gamma \leq \beta$ then this contradicts the choice of x_β, while for $\beta < \gamma$ this contradicts the choice of y_γ. Let W_I be any subset of A^ω such that $Y \subset W_I$ and $W_I \cap X = \emptyset$.
From the construction it follows that for each strategy σ of player I $H(\sigma)$ is not a subset of W_I and for each strategy τ of player II $H(\tau)$ is not a subset of $W_{II} = A^\omega \setminus W_I$, the players do not have winning strategies. □

1.2 Subgame perfect equilibria

Example 2. Figure 1 represents a perfect information extensive game with two players 1 and 2. There are two (pure) Nash equilibria (b, B) and (a, A). However (b, B) does not look very convincing. Suppose that player 1 plays rather a than b. When the play is at the position controlled by player 2 would he stick to B or rather would he play A? The threat of playing B does not seem to be convincing.

This example motivates the following definition.
A *subgame* of a game is defined as the game starting from some non terminal position x.

Definition 4. A strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a *subgame perfect equilibrium* in an extensive game Γ if σ is an equilibrium in each subgame of Γ.

Note that (b, B) is not a subgame perfect equilibrium in the game examined in Example 2. In any subgame perfect equilibrium player 2 will play A from the position that he controls.
For perfect information games we can find subgame perfect equilibria and payoffs of all players by backward induction. This is illustrated in Figure 3. Let us note that using backward induction we can always find subgame perfect equilibria in pure strategies. Since subgame perfect equilibria are also Nash equilibria we can see that extensive games with perfect information have Nash equilibria in pure strategies.

Example 3. Let us consider the following game (the centipede game).

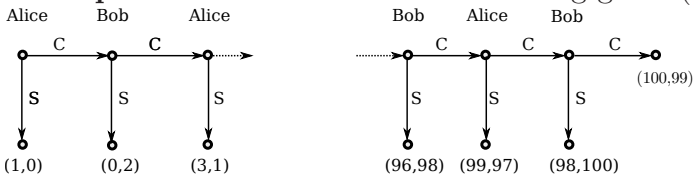

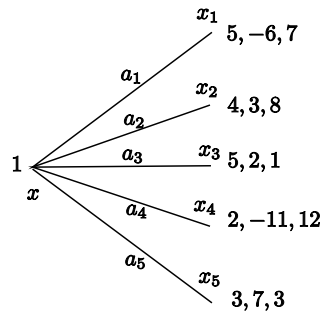

Figure 3: Suppose that x is a position controlled by player 1 in a perfect information game played by three players. Suppose that we have found subgame perfect equilibria in the subgames starting at x_1, x_2, x_3, x_4 and x_5 and that the three numbers labeling each of these positions give the payoffs of all three players in equilibrium. Clearly the optimal move of player 1 is to play either a_1 or a_3 or, more generally, he can use any randomization $\alpha a_1 + \beta a_3$, where $\alpha + \beta = 1$ and $\alpha, \beta \geq 0$. In equilibrium this will give him the payoff 5 while the actions a_2 and a_4 would give him only 4 and 2. If player 1 uses a strategy $\alpha a_1 + \beta a_3$ at x then the equilibrium payoff associated with x is $\alpha \cdot (5, -6, 7) + \beta \cdot (5, 2, 1) = (5, -6\alpha + 2\beta, 7\alpha + \beta)$.

For each terminal position the first number give the payoff of Alice and the second is the payoff of Bob.
At each position but the last one players choose between C (continue) and S (stop). The payoffs are arranged in such a way that each player prefers to stop immediately if the adversary stops in the next state. However if the adversary does not stop in the next state then the player prefers to continue.
What is the subgame perfect equilibrium? Does it seem reasonable? Would you play using the strategies of subgame perfect equilibrium?
Experiments show that human players never play a subgame perfect equilibrium in this game. □

Example 4. The bargaining game of alternating offers. Two players bargain over the division of a pie (we suppose that the pie is infinitely divisible). The play has a potentially unbounded number of rounds.
At first round player 1 offers to player 2 the share $0 \leq x_1 \leq 1$. If player 2 accepts (action A) then player 1 receives $1 - x_1$ and player 2 receives the proposed share x_2. If player 2 rejects (action R) then with probability $1 - \delta$ the game ends with payoff 0 for both players (the pie is declared improper for the consumption and thrown away). With probability δ we pass to the next round where player 2 makes an offer of x_2 to player 2. Either player 1 accepts and then receives the proposed x_2 while player 2 receives $1 - x_2$ or player 1 refuses. If player 1 refuses then with probability $1 - \delta$ the game ends with payoffs $(0, 0)$ and with probability δ we pass to the next round where player 1 will make some offer x_3. This game continues until one of the player accepts the offer proposed by the other player or the game stops by the chance move. Figure 4 presents a fragment of the game tree. Note that at a_1 we have in fact infinitely many outgoing actions, one action for each $x_1 \in [0, 1]$. The same remark holds for positions a_2 and a_3.

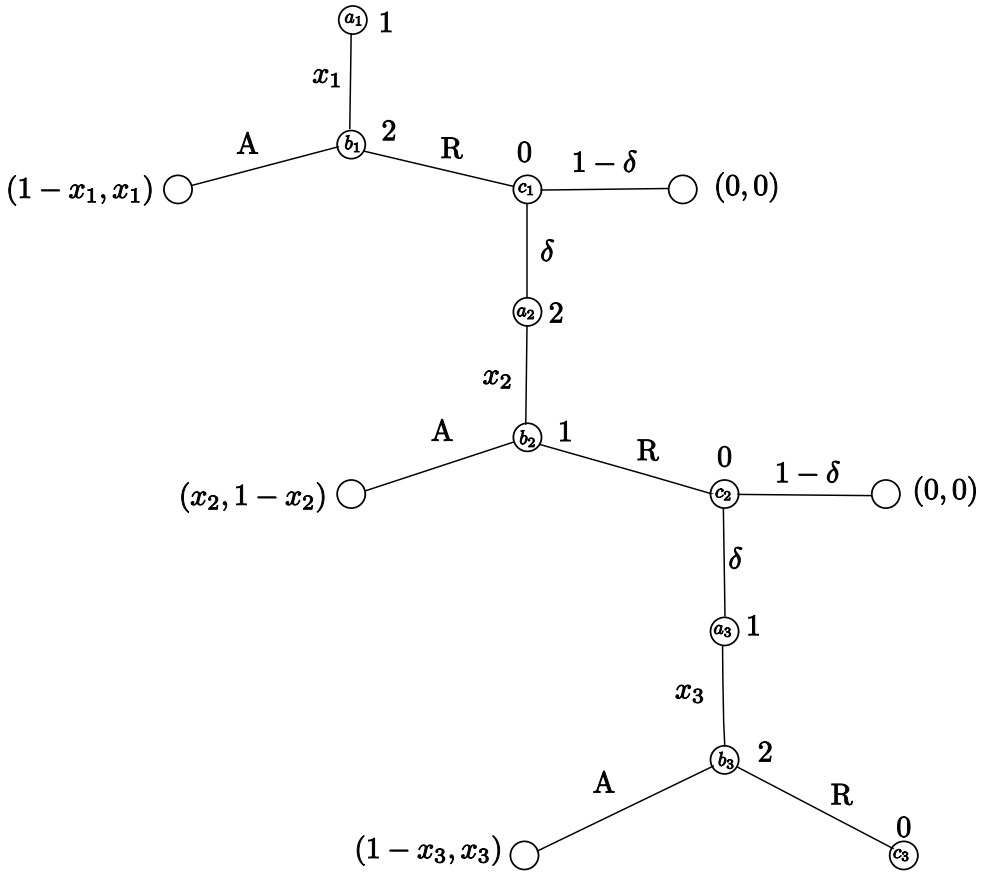

Figure 4: A part of the game tree in the bargaining game

For each $x \in [0,1]$ there is a Nash equilibrium with payoffs $(x, 1-x)$. Indeed strategies where player 1 never offers more than $1-x$ and never accepts less then x while player 2 never offers more then x and never accepts less than $1-x$ form a Nash equilibrium.
However we want to find subgame perfect equilibria.
Suppose that z is the maximum that player 1 can get in any subgame perfect equilibrium.

- Since the subtree rooted in a_3 is isomorphic to the whole tree, z is also the maximum that player 1 can get in the subgame rooted at a_3. Thus player 2 gets at least $1-z$ in the subgame rooted at a_3.

- Going backward, at c_2 player 1 wins at most $\delta z + (1-\delta)0 = \delta z$ while player 2 wins at least $\delta(1-z) + \delta 0 = (1-\delta)z$.

- This implies that player 2 will offer at a_2 at most δz to player 1 (there is no reason to offer more since δz is the maximum that player 1 can obtain by rejecting the offer).

- If player 1 accepts the offer above then player 1 gets δz while player 2 gets $1-\delta z$. Thus at a_2 player 1 gets at most δz while player 2 gets at least $1-\delta z$.

- This implies that at c_1 player 1 gets at most $\delta(\delta z)+(1-\delta)0 = \delta^2 z$ while player 2 gets at least $\delta(1-\delta z)+(1-\delta)0 = \delta - \delta^2 z$.
- However, this implies that at a_1 player 1 will not offer more than $\delta - \delta^2 z$ (no reason to offer more since player 2 cannot win more by rejecting this offer).
- Since player 2 is compelled to accept this generous offer, at a_1 player 1 wins at most $1-(\delta - \delta^2 z)$ while player 2 wins at least $\delta - \delta^2 z$.

On the other hand player 1 wins at most z thus we have $1-(\delta - \delta^2 z) = z$, implying and $z = 1/(1+\delta)$.
We can repeat a similar analysis supposing that w is the minimum that player 1 can obtain in any subgame perfect equilibrium. This would yield the same equation as for z. Thus $1/(1+\delta)$ is what player 1 wins in any subgame perfect equilibrium. We conclude that

$$(\frac{1}{1+\delta}, \frac{\delta}{1+\delta})$$

is the payoff of both players in the subgame perfect equilibrium.
Instead of supposing that the game ends with probability $1-\delta$ at the end of each round in the case of refusal we can assume that the pie shrinks by factor δ at the beginning of each round, i.e. at the first round the players have a pie of size 1, at the second round a pie of size δ, at the third round a pie of size δ^2 etc. and after each refusal we pass directly to the next round, there is no chance move that can stop the game. The game analysis and the result are the same for this variant[1].

Exercise 1 (From Herbert Gintis, Game theory evolving)**.** Miss Muffet is eating an ice cream and is confronted with a wasp which apparently likes to take its share.
Miss Muffet proposes the following bargaining game. She proposes a certain share x. If the wasp accepts the offer then each player goes away with her/its share. If the offer is rejected then in the next round the wasp makes a counter offer under the same conditions as above. However the day is hot, by the time the second offer is accepted or rejected the ice have melted down to the half of the original size. Both Miss Muffet and the wasp know perfectly all these conditions, are fully rational and they refuse the offer only if they can gain something in this way. We assume that after the second rejection Miss Muffet and the wasp get 0 and there is no third round in this game, by this time the ice cream will melt down completely.
How much Miss Muffet offers at the first round? Is this offer accepted?

For the sequel suppose that the whole ice cream is of size 1.
Consider the case where there are even number n of rounds, Miss Muffet and the wasp make the offers alternately, with Miss Muffet making the first offer, and after each refusal the ice cream shrinks by $\frac{1}{n}$, i.e. after k refusal the size of the ice cream at the beginning of the $(k+1)$st round is $\frac{n-k}{n}$ (in particular after n refusals the ice cream melts down completely). Show that Miss Muffet will offer $\frac{1}{2}$ of the ice cream and this offer will be immediately accepted.
Now suppose that with the same conditions as above the number of rounds n is odd. Show that Miss Muffet will offer $\frac{1}{2} - \frac{1}{2n}$ and this offer will be accepted.

[1]Since the chance cannot stop this variant of the bargaining game there exist strategy profiles with infinite plays, for such plays we assume the payoffs $(0,0)$.

Exercise 2. *Warning: This is a rather difficult exercise, still you should try to solve it.* Jacques, Pierre and Magali want to split a pie. They proceed as in Example 4, i.e. at first round Jacques proposes to split the pie giving p to Pierre, m to Magali and keeping j to himself. If Pierre and Magali accept then the pie is shared as proposed. If either Pierre or Magali refuse then with probability $1-\delta$ the game stops and with probability δ we pass to the next round where Pierre makes a proposal. If this proposal is rejected either by Magali or by Jacques then in the next round it is Magali's turn to make a proposal. The players turn this way unless the game is stopped by the nature or a proposal of one player is accepted by the other two players. Show that there is a unique symmetric and stationary[2] subgame perfect equilibrium.
There exist also non symmetric subgame perfect equilibria, but they are difficult to find and to describe.

2 Extensive games with imperfect information

An extensive game with imperfect information has the same ingredients as a game with perfect information but now we assume that each set V_i, $i \in N$, is partitioned onto k_i non empty sets, $V_i = V_i^1 \cup \dots V_i^{k_i}$, the so called *information sets*.
We assume that for any two positions v, w that are in the same information set we have $A(v) = A(w)$, i.e. v and w have the same sets of available actions. In the sequel, for each information set V_i^m, $A(V_i^m)$ will denote the set of actions available at V_i^m.
The intuition behind the information sets is that if the play arrives at any position belonging to information set V_i^j then player i knows only that he is at some position in V_i^j but he does not know which one.
If all information sets are singletons then we have a perfect information game examined in the preceding section.
An example of an extensive game with imperfect information is given on Figure 5.

Example 5. Figure 5 presents a game with partial information. At state r there is

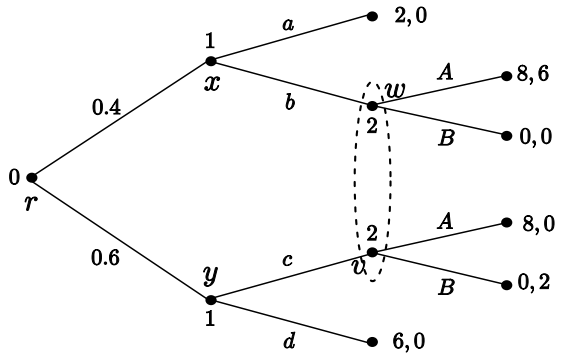

Figure 5: A extensive game with imperfect information.

a chance move, with probability 0.4 we go to the state x and with probability 0.6

[2]Symmetric stationary means that all players make the same offer at each round.

to state y. Both these state belong to player 1, actions a and d lead to terminal states and actions b and c lead to states w and v controlled by player 2. Player 2 cannot distinguish between these two states, he can observe neither chance moves nor player's 1 moves.

2.1 Pure, mixed and behavioral strategies

Let Γ be a game in extensive form. Let $V_i^1, \ldots, V_i^{k_i}$ be information sets of player i.
A *pure strategy* for player i is a mapping σ_i that for each information set V_i^j gives an action $\sigma_i(V_i^j) \in A(V_i^j)$ available in V_i^j.
A *mixed strategy* for player i is a probability distribution over pure strategies of player i.
A *behavioral strategy* σ_i for player i is a mapping that for each information set V_i^j gives a probability distribution over the actions available at V_i^j, $\sigma_i(V_i^j) \in \Delta(A(V_i^j))$.

Example 6. Let us consider the game in Figure 6.

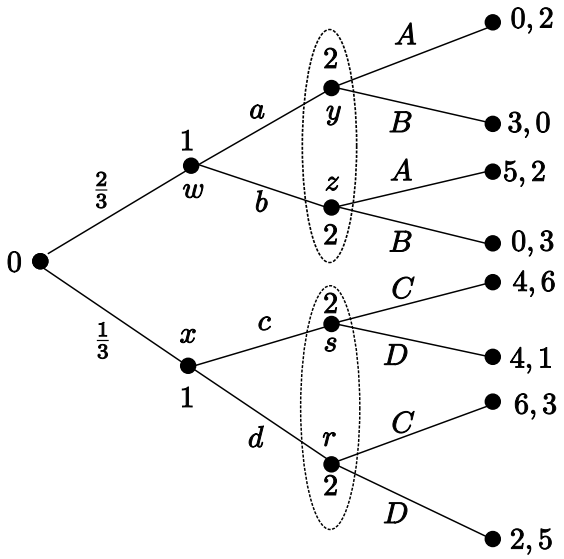

Figure 6: An extensive game with imperfect information. There is one chance position – the root. The information sets of player 1 are singletons. There are two information sets for player 2.

Player 1 has two information sets consisting of single vertices and he has 4 pure strategies: $[ac], [ad], [bc], [bd]$. Player 2 has also two information sets and four pure strategies: $[AC], [AD], [BC], [BD]$. Any profile of pure strategies yields a probability distribution over terminal positions, for example if players use strategies $([ac], [AD])$ then terminal vertices with payoffs $(0, 2)$ and $4, 1$ are obtained with probabilities $\frac{2}{3}$ and $\frac{1}{3}$ respectively, all other terminal positions have probability 0. Thus we can associate with profile $([ac], [AD])$ the payoffs $\frac{2}{3}(0, 2) + \frac{1}{3}(4, 1) = (\frac{4}{3}, \frac{5}{3})$. The same method can applied to any pure strategy profile to calculate the corresponding payoffs.
This gives a translation from extensive games to normal form games and the Nash theorem implies that the existence of equilibria in mixed strategies.

However we are more interested in equilibria in behavioral strategies.

Example 7. Consider the game presented in Figure 7.

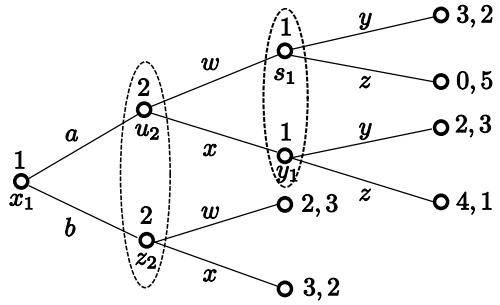

Figure 7: An extensive game with imperfect information. There are two information sets for player 1, $\{x_1\}$ and $\{s_1; y_1\}$, and one information set $\{u_2, z_2\}$ of for player 2.

An example of a behavioral strategy for player 1 is $\sigma_1(x_1) = \frac{2}{5}a + \frac{3}{5}b$ and $\sigma_1(\{s_1, y_1\}) = \frac{1}{4}y + \frac{3}{4}z$, i.e. in the information set $\{x_1\}$ actions a and b are taken with probability $\frac{2}{5}$ and $\frac{3}{5}$ respectively and in the information set $\{s_1, y_1\}$ player 1 takes action y with probability $\frac{1}{4}$ and action z with probability $\frac{3}{4}$.
Again each profile of behavioral strategies yields a probability distribution over terminal vertices and taking the expectation we get the payoff of each player. But does there exist an equilibrium in behavioral strategies? The answer is positive for games with perfect recall.

2.2 Games with perfect recall

Let Γ be an extensive game (with imperfect information). For each path $p = v_0 a_0 v_1 \dots a_{n-1} v_n$ from the initial position v_0 to a position v_n we can define the view $\text{view}_i(p)$ of player i in the following way:

- first remove from p all states that are not in V_i and all actions that are not executed by player i,
- in the next step replace the states of V_i by the corresponding information sets.

Definition 5. A game is a *game with perfect recall* iff for each player i, for each information set V_i^j of player i and for all positions $v, w \in V_i^j$ if p_v and p_w are paths from the initial position to v and w respectively then $\text{view}_i(p_v) = \text{view}_i(p_w)$. In other words, player i cannot distinguish v from w even if he has a perfect memory of all events that he could observe during the play.

The games in Figures 7, 6 and 5 are with perfect recall.

Example 8. Figure 8 provides an example of a game without perfect recall. Let p_z and p_y be the paths from the initial position to z and y respectively. We have $\text{view}_1(p_z) = \{z, y\}$ and $\text{view}_1(p_y) = \{x\}c\{z, y\}$. Thus if player 2 confuses z and y then this means that he "has forgotten" whether he had visited x or not and whether he had taken action c. □

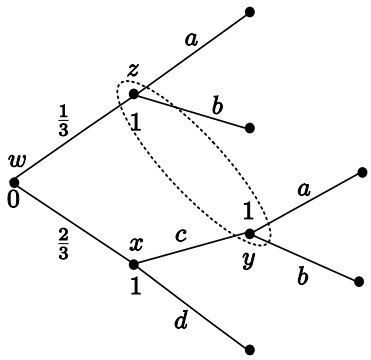

Figure 8: There are two information sets for player 1, $\{x\}$ and $\{z, y\}$.

Example 9 (Absent-minded driver)**.**

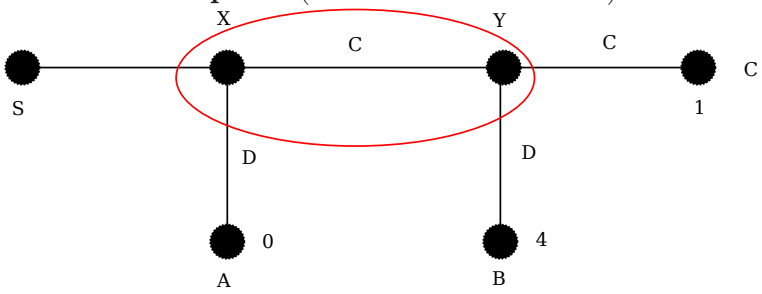

A driver takes a highway at S. His aim is to go to his home at B (the highest payoff). However to go to B he should exit the highway at Y. If he exits the highway earlier at X he will arrive at some very nasty place A (the lowest payoff). If he never exits he will arrive at C when he is obliged to take a hotel room (better then A but definitely worst than his home).
However he is absent-minded and he confounds exits X and Y. □

Notation: For a strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ and a player $i \in N$, by $\sigma_{-i} = (\sigma_1, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_n)$ we denote the strategy profile of all players exept player i.
On the other hand if $\sigma_{-i} = (\sigma_1, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_n)$ is a strategy profile of all players except i and if τ_i is a strategy of i then (τ_i, σ_{-i}) denotes the strategy profile $(\sigma_1, \ldots, \sigma_{i-1}, \tau_i, \sigma_{i+1}, \ldots, \sigma_n)$.

Definition 6. Two strategies σ_i and σ_i' (mixed or behavioral) of player i are said to be *outcome-equivalent* if for any strategy profile σ_{-i} of the other players and for each terminal position $t \in T$, $P(t|(\sigma_i, \sigma_{-i})) = P(t|(\sigma_i', \sigma_{-i}))$, i.e. the probability of reaching t under (σ_i, σ_{-i}) is the same as the probability of reaching t under (σ_i', σ_{-i}).

Theorem 7 (Kuhn)**.** *For finite games with perfect recall mixed and behavioral strategies are outcome-equivalent.*

Corollary 8. *Finite games with perfect recall have Nash equilibria in behavioral strategies.*

We just sketch briefly the proof of Theorem 7.

The translation from behavioral to outcome-equivalent mixed strategies is easy. Consider again the game of Figure 7. The behavioral strategy

$$\{x_1\} \mapsto \alpha a + (1-\alpha)b, \qquad \{s_1, y_1\} \mapsto \beta y + (1-\beta)z$$

is outcome-equivalent to the mixed strategy

$$\alpha\beta[ay] + \alpha(1-\beta)[az] + (1-\alpha)\beta[b,y] + (1-\alpha)(1-\beta)[yz].$$

I hope that you get the pattern: If $V_i^1, \ldots, V_i^{i_k}$ are information sets of player i, $A_i^1, \ldots, A_i^{i_k}$ are the sets of actions available at each information set and σ a behavioral strategy that maps each V_i^j to an element of $\Delta(A_i^j)$ then the corresponding outcome-equivalent mixed strategy σ^* is obtained in the following way: the probability of a pure strategy $[a_1, \ldots, a_{i_k}]$, where $a_1 \in A_1, \ldots, a_{i_k} \in A_i^{i_k}$ is simply $\sigma^*([a_1, \ldots, a_{i_k}]) = \sigma(V_1)(a_1) \cdots \sigma(V_i^{i_k})(a_{i_k})$.
The proof that σ^* is outcome-equivalent to σ is easy and can be left as an exercise.

The translation from a mixed to an outcome-equivalent behavioral strategy needs more care as the following example shows:
Consider the mixed strategy $\frac{1}{2}[ay]+\frac{1}{2}[bz]$ of player 1 in the game of Figure 7. What is the equivalent behavioral strategy? In particular what is the probability distribution for actions y and z at the information set $\{s_1, y_1\}$? One can think naively that y, z are executed with respective probabilities 0.5 and 0.5. Note however that player 1 using the mixed strategy $\frac{1}{2}[ay] + \frac{1}{2}[bz]$ will never play z, this is because if he chooses b at x_1 then the information set $\{s_1, y_1\}$ is never attained, thus the probability of choosing z in an equivalent behavioral strategy is 0 and not 0.5.

To construct a behavioral strategy equivalent to a given mixed strategy we proceed in the following way.
A *history* is a path $h = s_1 a_1 s_2 a_2 \ldots a_{n-1} s_n$ in the game tree starting at the root s_1 of the game tree and ending at some position s_n.
Let β be a pure strategy for player i. We say that $h = s_1 a_1 s_2 a_2 \ldots a_{n-1} s_n$ is *consistent* with β if for each position s_m in h which is controlled by player i, $\beta(V_i^{s_m}) = a_m$, where $V_i^{s_m}$ is the information set containing s_m.
Let h and h' be two histories ending at vertices (positions) which are in the same information set V_i^j of player i. The crucial observation is that for a perfect recall game and any pure strategy β of player i, h is consistent with β if and only if h' is consistent with β (and this property does not hold for games without perfect recall).
Let B_i be the set of all pure strategies of player i.
For each information set V_i^j of player i, let $B_i(V_i^j)$ be the set of all pure strategies $\beta \in B_i$ of player i such that there exists a history h ending at V_i^j which is consistent with β (by our previous remark concerning perfect recall games, this happens if all histories ending at V_i^j are consistent with β).
For an action $a \in A(V_i^j)$ available at V_i^j, let $B_i(V_i^j, a)$ be the set of all pure strategies $\beta \in B_i$ of player i such that $\beta \in B_i(V_i^j)$ and $\beta(V_i^j) = a$.
Let $\sigma = \sum_{\{\beta \in B_i\}} b_\beta \cdot \beta$ be a mixed strategy of player i, where b_β is the probability of choosing pure strategy β, $\sum b_\beta = 1$.
Then a behavioral strategy σ^* outcome equivalent to σ is defined in the following way:

for each information set V_i^j of player i and each action $a \in A(V_i^j)$ we set

$$\sigma^*(V_i^j)(a) = \frac{\sum_{\beta \in B_i(V_i^j, a)} b_\beta}{\sum_{\beta \in B_i(V_i^j)} b_\beta} \tag{2}$$

if $\sum_{\beta \in B_i(V_i^j)} b_\beta \neq 0$. If $\sum_{\beta \in B_i(V_i^j)} b_\beta = 0$ then $\sigma^*(V_i^j)(a)$ can be set to any constant from the interval $[0, 1]$ (in such a way that the the sum of $\sum_{a \in A(V_i^j)} \sigma^*(V_i^j)(a) = 1$). Again the proof that σ and σ^* are outcome-equivalent is left as an exercise.
The intuitive meaning of (2) is the following. Denominator gives a probability that the mixed strategy σ select a pure strategy that can lead to the information set V_i^j while the numerator is the probability that σ select a pure strategy thac can lead to V_i^j and that will choose action a in V_i^j.

Example 10. Let us consider the game of Example 7. Player 1 has four pure strategies $[ay]$ $[az]$, $[by]$ and $[bz]$. Let

$$\sigma = b_1[ay] + b_2[az] + b_3[by] + b_4[bz]$$

be a mixed strategy of player 1.
Take the information set $\{s_1, y_1\}$ of player 1. Then $B_1(\{s_1, y_1\}) = \{[ay], [az]\}$, i.e. histories ending at $\{s_1, y_1\}$ are consistent either with $[ay]$ or with $[az]$, and they are consistent neither with $[by]$ nor with $[bz]$. We have also $B_1(\{s_1, y_1\}, y) = \{[ay]\}$, i.e. $[ay]$ is the only pure strategy such that histories ending at $\{s_1, y_1\}$ are consistent with it and this strategy selects action y. Thus a behavioral strategy σ^* outcome equivalent to σ selects action y at $\{s_1, y_1\}$ with probability

$$\sigma^*(\{s_1, y_1\})(y) = \frac{b_1}{b_1 + b_2}$$

if $b_1 + b_2 \neq 0$. If $b_1 + b_2 = 0$ then we can set $\sigma^*(\{s_1, y_1\})(y)$ anyhow. □

2.3 Subgame perfect equilibria for extensive games with imperfect information

Subgame perfect equilibria for imperfect information games are defined in the same way as for perfect information games if we adopt the following definition of the subgame: a subgame is a subtree starting at some non-terminal position and such that that for each player i and for each information set V_i^j of player i either V_i^j is contained in the subtree or it is disjoint with the subtree. For example the game in Figure 7 has no proper subgames. On the other hand, the game of Figure 6 has two proper subgames rooted at w and x.
There many other notions of equilibria for imperfect information games: sequential equilibria, trembling hand equilibria etc.

Exercise 3 (from David Williams, Weighing the odds, Cambridge University Press)**.** At the end of Monty Hall game show a contestant is shown three closed doors. Behind one of the doors is a car; behind each of the other two is a goat. The contestant chooses one of the three doors. The show's host, **who knows which**

door conceals the car, opens one of the remaining two doors which he knows will definitely reveal a goat. He then asks the contestant whether or not she wishes to switch her choice to the remaining closed door. Should she switch or stick to the original choice?
Model this as a game with imperfect information. Answer the question given above. Note that there is a number of hidden assumptions. First of all we suppose that the contestant prefers a car to a goat (maybe this is obvious?).
Secondly, all this description is common knowledge, in particular the contestant knows that the host knows where is the car and that the host will always open the door with a goat behind.
The problem can be solved by elementary probability but I want that you model it as a game and solve on the basis of the game presentation.
Now suppose that the host does not know where is the car, he opens the door and it turns out that there is a goat behind. In this case should the contestant switch or stick to the original choice?

3 Combinatorial games

The nim game is played with a (finite) number of piles of chips. A legal move is to strictly decrease the number of chips in any pile (and throw away the removed chips).
The nim-sum $a \oplus b$ of two natural numbers a and b is obtained by writing them down in binary and adding them without carrying[3] .
Example:

$$5 \oplus 9 = (101)_2 \oplus (1001)_2 = (1100)_2 = 12$$

Obviously $(\mathbb{N}, \oplus)$ is an abelian group.

Theorem 9 (Bouton 1902)**.** *Suppose that we have p piles of size $n_1, \ldots, n_p$ respectively. Then position $(n_1, \ldots, n_p)$ is winning for the current player in the nim game iff $n_1 \oplus \ldots \oplus n_p \neq 0$.*

Proof. Let

$$V = \{(n_1, \ldots, n_p) \mid \forall i, n_i \geq 0\}$$

be the set of all positions,

$$W = \{(n_1, \ldots, n_p) \mid n_1 \oplus \ldots \oplus n_p \neq 0\}$$

the set of winning positions (for the current player),

$$L = \{(n_1, \ldots, n_p) \mid n_1 \oplus \ldots \oplus n_p = 0\}$$

the set of losing positions (losing for the current player).
Note that the only terminal position $(0, \ldots, 0)$ in the nim game, with all piles empty, belongs to L. Therefore to prove that the elements of W and L are effectively winning and losing positions respectively it is sufficient to show that

[3]In other words, this is just bit-wise XOR of binary representations of a and b.

(A) for each position in W there exists a move to a position in L (the player playing in such a position can force the game to a losing position),

(B) for each non-terminal position in L each available move leads to W (player playing in such a position cannot avoid a "bad" move). .

The second assertion is easy.
Let

$$v = (n_1, \ldots, n_{i-1}, n_i, n_{i+1}, \ldots, n_p) \in L,$$

and

$$w = (n_1, \ldots, n_{i-1}, m_i, n_{i+1}, \ldots, n_p)$$

with $n_i > m_i \geq 0$, i.e. (v, w) is a valid move from a losing position v decreasing the i-th pile.
Let us note

$$k = n_1 \oplus \ldots \oplus n_{i-1} \oplus n_{i+1} \oplus \ldots \oplus \oplus n_p$$

the nim-sum of all piles except the i-th one.
Then, by definition of L,

$$k \oplus n_i = 0,$$

which is possible if and only if $k = n_i$. But then $k \neq m_i$ and

$$n_1 \oplus \ldots \oplus n_{i-1} \oplus m_i \oplus n_{i+1} \oplus \ldots n_p = k \oplus m_i \neq 0$$

i.e. $w \in W$. This ends the proof of (B).
To prove (A) let us suppose that

$$w = (n_1, \ldots, n_p) \in W$$

i.e.

$$n_1 \oplus \ldots \oplus n_p \neq 0$$

Let

$$a = n_1 \oplus \ldots \oplus n_p.$$

For a non-negative integer x we shall note by $(x)_i$ the i-th bit in the binary development of x.
Let $l = \max\{i \mid (a)_i = 1\}$ be the most significant bit of a equal to 1.
Since $(a)_l = 1$ at least one of the numbers n_i has the l-th bit 1. This implies that $n_i \oplus a < n_i$ since

- the numbers n_i and $n_i \oplus a$ have the same bits on positions greater l-th (because all bits of a at positions higher than l-th are 0),
- the l-th bit of $n_i \oplus a$ is 0 $(1 \oplus 1 = 0)$
- while the l-th bit of n_i is 1.

Therefore the action that the player should take at v consists in reducing the size n_i of the i-th pile to $m_i = n_i \oplus a$.
Note that

$$n_1 \oplus \ldots \oplus m_i \oplus \ldots \oplus n_p = n_1 \oplus \ldots \oplus (n_i \oplus a) \oplus \ldots \oplus n_p = \\ n_1 \oplus \ldots \oplus n_i \oplus \ldots \oplus n_p \oplus a = a \oplus a = 0$$

i.e. the new position $(n_1, \ldots, m_i, \ldots, n_p)$ attained by this move belongs to L. □

4 Examples

Several games can be reduced to nim.

Poker nim

As the nim game this game is played with piles of chips. But now each player has also a private stack of chips (initially the private stack of each player can contain some finite number of chips).
There are two possible valid moves:

- the player may remove any (positive) number of chips from one of the piles like in the nim game or
- the player the can add any (positive) number of chips from his private stack to one of the piles.

This is not a combinatorial game since the corresponding game graph can have cycles.
However, it turns out that one of the players has a winning strategy assuring that the game ends after a finite number of steps with the other player unable to move (i.e. at the end it is the other player turn to move but all piles as well as the private stack of the other player are empty).
This game can be reduced to the nim game, in fact the winning and losing positions are determined exactly in the same way as in as in the nim game (and the chips in the private stacks do not count when we determine winning/losing positions!).
So the only question that the winning player should answer is what to do if the adversary tries to distract him by increasing the number of chips in some pile, what is the appropriate counter-move in this case? (**Answer this question as an exercise!**).

Northcott's game

Played on the checkerboard. One white and one black token in each row. You may move a token of your color to any other empty square in the same row provided that you do not jump over your opponent's token.

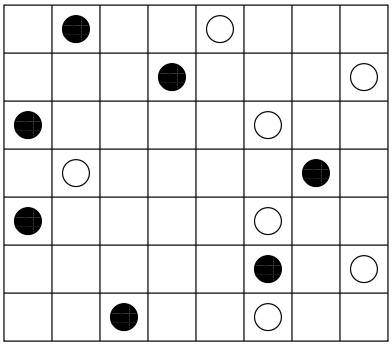

If you cannot move you lose. Again this game can be easily reduced to the nim game (**exercise**).

A silver dollar game

Played on a right-infinite tape with a finite number of coins placed on squares (each coin on a separate square).

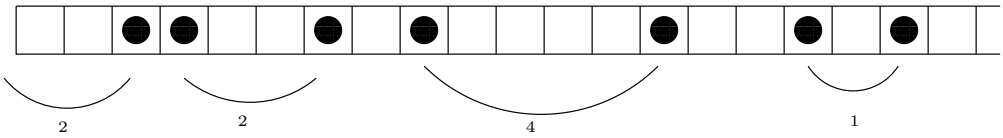

A legal move is to move a coin leftwards on a free square (there should be always at most one coin at a square) without passing over any other coin. The game ends when there is no legal move, i.e. when all coins are in a traffic jam at the left end of the tape.

Again this is the nim game in disguise. Starting from the rightmost coin calculate the number of squares in alternate spaces between coins, take this as the size of the nim piles. In the picture above we have 4 nim piles of sizes $2, 2, 4, 1$.

A position is winning if and only if the corresponding configuration of nim piles is winning.

Exercise Find the strategy for the winning player (again the crucial point is what to do if the adversary increases the size of a pile by moving the coin on the left border of the pile).